

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE  
APPLICATION FOR LETTERS PATENT

**Controlling Memory Usage in Systems Having Limited  
Physical Memory**

Inventor(s):

Sharad Mathur  
Michael Ginsberg  
Thomas Fenwick  
Anthony Kitowicz  
William H. Mitchell  
Jason Fuller

1            **TECHNICAL FIELD**

2            This invention relates to computer systems having virtual memory systems  
3            that utilize limited physical memory and no secondary storage.

5            **BACKGROUND OF THE INVENTION**

6            Most modern operating systems implement virtual memory. Virtual  
7            memory is a seemingly large amount of memory that allows individual application  
8            programs in a multi-tasking system to use respective, dedicated address spaces.  
9            Each dedicated address space includes addresses from zero to some large number  
10          that depends on the particular characteristics of the operating system and the  
11          underlying hardware.

12          In a virtual memory system, an application program is assigned its own  
13          virtual address space, which is not available to other application programs.  
14          Through its virtual memory, a process has a logical view of memory that does not  
15          correspond to the actual layout of physical memory. Each time a program uses a  
16          virtual memory address, the virtual memory system translates it into a physical  
17          address using a virtual-to-physical address mapping contained in some type of  
18          look-up structure and address mapping database.

19          Rather than attempting to maintain a translation or mapping for each  
20          possible virtual address, virtual memory systems divide virtual and physical  
21          memory into blocks. In many systems, these blocks are fixed in size and referred  
22          to as sections or pages. Data structures are typically maintained in physical  
23          memory to translate from virtual page addresses to physical page addresses.  
24          These data structures often take the form of conversion tables, normally referred to  
25          as page tables. A page table is indexed by virtual page address or number, and

1 generally has a number of entries corresponding to pages in the virtual address  
2 space. Each entry is a mapping of a specific page number or virtual page address  
3 to a physical page address.

4 In most virtual memory systems, physical memory includes some form of  
5 secondary storage such as a hard disk. When primary, electronic memory  
6 becomes full, physical memory pages are moved to the disk until they are accessed  
7 again. This process is referred to as paging. Assuming that the hard disk has a  
8 large capacity, paging allows the simulation of seemingly unlimited physical  
9 memory.

10 Although hard disks are common, and are becoming less and less  
11 expensive, there is a new generation of computers that implement virtual memory  
12 systems without the benefit of secondary storage. Currently, these computers  
13 primarily comprise so-called handheld computers or "H/PCs" (handheld PCs).  
14 H/PCs typically have a limited amount of non-volatile addressable memory such  
15 as battery-backed dynamic RAM (random access memory). Some of this memory  
16 is allocated for program execution, while the remaining memory is used to  
17 implement a file system. While more capable H/PCs might include actual hard  
18 disk storage, this is not the usual situation.

19 These computers impose new restrictions on the use of virtual memory. In  
20 systems that included secondary storage, there was little danger of exhausting  
21 physical memory since it could be paged to disk. In an H/PC, however, care must  
22 be taken to conserve memory usage. In a multi-tasking system, it is possible to  
23 launch a program that competes with other programs and with the operating  
24 system for available memory. If any particular program makes high memory  
25 demands, it is conceivable that other programs might find themselves without

1 enough memory to continue. Even worse, it is possible that the operating system  
2 itself could be unable to obtain needed memory, thereby causing a system crash.

3 It would be desirable to limit memory usage only when required, rather  
4 than prospectively limiting application programs to prescribed memory usage  
5 limits. However, it would also be desirable to prevent application programs from  
6 threatening system stability. The system and methods described below  
7 accomplish these goals.

8

## 9 SUMMARY OF THE INVENTION

10 The invention is implemented within an operating system that continually  
11 or periodically monitors memory usage. Three usage thresholds are established,  
12 and different actions are taken as increasingly critical memory usage thresholds  
13 are reached. At each threshold, the objective is to free memory so that the next  
14 higher threshold is avoided.

15 At the first, least critical threshold, one or more application programs are  
16 simply requested to minimize their memory usage. The request is issued to the  
17 least recently active programs, through their Windows® message loops. The  
18 applications can respond to the requests as they deem appropriate. Well-behaved  
19 applications will take steps to release resources as much as possible.

20 At the second, more critical threshold, the operating system closes one or  
21 more of the application programs. The application is closed using a standard  
22 operating system mechanism that allows the application to shut down in an orderly  
23 fashion, while saving files and performing any other necessary housekeeping.

24 At the third, most critical threshold, the operating system simply terminates  
25 one or more of the application programs. The application program's thread(s) are

1 destroyed and all resources previously used by the application program are  
2 released. The application program is given no opportunity to clean up, to save  
3 files, or to perform any other housekeeping chores.

4 Other measures are potentially employed before taking any of the three  
5 measures described above. For example, an attempt is made to reclaim any  
6 unused stack space. In addition, read-only memory pages are discarded. Such  
7 preventative measures are preferably initiated at further memory usage thresholds.  
8 These further thresholds preferably have a fixed relationship to the first, second,  
9 and third thresholds noted above.

10

11 **BRIEF DESCRIPTION OF THE DRAWINGS**

12 Fig. 1 is shows a handheld computing device 20 in accordance with one  
13 embodiment of the invention.

14 Fig. 2 is a block diagram of the device shown in Fig. 1.

15 Fig. 3 is a simplified flow chart showing major steps in accordance with the  
16 invention.

17

18 **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT**

19 Fig. 1 shows a handheld computing device 20 such as might be utilized to  
20 implement the invention. As used herein, "handheld computing device" means a  
21 small computing device having a processing unit that is capable of running one or  
22 more application programs, a display, and an input mechanism that is typically  
23 something other than a full-size keyboard. The input mechanism might be a  
24 keypad, a touch-sensitive screen, a track ball, a touch-sensitive pad, a miniaturized  
25 QWERTY keyboard, or the like.

1       The handheld computing device 20 of Fig. 1 is also referred to as a  
2 handheld personal computer (H/PC). However, in other implementations, the  
3 handheld computing device may be implemented as a personal digital assistant  
4 (PDA), a personal organizer, a palmtop computer, a computerized notepad, or the  
5 like. The invention can also be implemented in other types of computers and  
6 computer-like or computer-controlled devices.

7       H/PC 20 has a casing 22 with a cover or lid 24 and a base 26. The H/PC  
8 has a liquid crystal display (LCD) 28 with a touch-sensitive screen mounted in lid  
9 24. The lid 24 is hinged to base 26 to pivot between an open position, which  
10 exposes screen 28, and a closed position, which protects the screen. The device is  
11 equipped with a stylus 30 to enter data through the touchscreen display 28 and a  
12 miniature QWERTY keyboard 32. Both of these components are mounted in base  
13 26. H/PC 20 can also be implemented with a wireless transceiver (not shown)  
14 such as an IR (infrared) transceiver and/or an RF (radio frequency) transceiver.  
15 Although the illustrated implementation shows a two-member H/PC 20 with a lid  
16 24 and a base 26, other implementations of the H/PC might comprise an integrated  
17 body without hinged components, as is the case with computerized notepads (e.g.,  
18 Newton® from Apple Computers).

19      Fig. 2 shows functional components of the handheld computing device 20.  
20 It has a processor 40, addressable memory 42, a display 28, and a keyboard 32. At  
21 least a portion of memory 42 is non-volatile. The memory is divided into a  
22 physical memory portion that is addressed by processor 40 for program execution,  
23 and a file portion that is used to implement a file system.

24      An operating system 44, executed by processor 40, is resident in and  
25 utilizes memory 42. H/PC 20 preferably runs the Windows® CE operating system

1 from Microsoft Corporation. This operating system is a derivative of Windows®  
2 brand operating systems, such as Windows® 95, that is especially designed for  
3 handheld computing devices. The invention is implemented within the  
4 Windows® CE operating system—the operating system includes instructions that  
5 are executable by processor 40 to implement the features and to perform the  
6 specific steps described below.

7 The Windows® CE operating system is a multitasking operating system  
8 that allows simultaneous execution of multiple applications 45. The operating  
9 system employs a graphical user interface windowing environment that presents  
10 applications and documents in specially delineated areas of the display screen  
11 called “windows.” Each window can act independently, including its own menu,  
12 toolbar, pointers, and other controls, as if it were a virtual display device. It is  
13 noted, however, that the handheld computing device may be implemented with  
14 other types of operating systems.

15 The H/PC 20 has a power supply 46 that supplies power to the electronic  
16 components. The power supply 46 is preferably implemented as one or more  
17 batteries. The power supply 46 might further represent an external power source  
18 that overrides or recharges the built-in batteries, such as an AC adapter or a  
19 powered docking cradle.

20 The Windows® CE operating system provides many of the same  
21 capabilities of more sophisticated operating systems such as Microsoft Windows®  
22 95 and Microsoft Windows® NT. Particularly, the Window® CE operating  
23 system implements a virtual memory system 48 that utilizes portions of physical  
24 memory 42. Thus, each application program 45 utilizes physical memory 42  
25 through virtual memory system 48. A notable characteristic of H/PC 20, however,

1       is the absence of secondary storage such as a hard disk. The virtual memory  
2       system in Windows® CE is limited to the available physical memory 42, and does  
3       not include secondary storage.

4       Operating system 44 implements a method of controlling memory usage in  
5       a system such as shown in Figs. 1 and 2, in which one or more application  
6       programs execute concurrently and in which the application programs compete  
7       with each other and with the operating system for available memory.

8       Generally, operating system 44 monitors memory usage and sets a plurality  
9       of memory usage thresholds at which different actions are taken to reduce or  
10      minimize current and future memory usage. In the actual embodiment of the  
11      invention, the memory thresholds are set in terms of remaining available or  
12      uncommitted memory pages, although the thresholds could also be specified in  
13      terms of used or committed memory.

14      Upon reaching increasingly critical memory thresholds, the operating  
15      system wields increasing control over one or more of application programs 45 to  
16      minimize memory usage. At less critical thresholds, the operating system politely  
17      interacts with application programs to limit or minimize their use of memory. For  
18      example, at a first or least critical memory usage threshold, the operating system  
19      requests one or more application programs to limit or minimize their current use of  
20      memory. At a second, more critical memory usage threshold, the operating system  
21      uses normal operating system messages to close application programs, thus  
22      allowing the application programs to shut down themselves.

23      At a third critical memory usage threshold that is more critical than the first  
24      two thresholds, the operating system simply and abruptly terminates one of the  
25      application programs, without allowing any further execution of the application

1 program. The application program is not given the opportunity to shut down  
2 itself.

3 Fig. 3 shows steps performed by operating system 44 in accordance with  
4 the invention. These steps are performed periodically or continuously, as  
5 described in more detail below, and do not necessarily follow the sequence  
6 suggested by the flowchart of Fig. 3. However, the flowchart is useful for  
7 understanding the relationships of the three memory usage thresholds used in the  
8 system.

9 A first step 100 comprises comparing current memory usage or availability  
10 against a usage or availability threshold that is deemed to be “most” critical. This  
11 memory threshold represents the highest level of memory usage, or the lowest  
12 level of free memory, that is considered acceptable for stable system operation. In  
13 Windows® CE, this is referred to as the critical memory threshold. If this  
14 threshold has been reached, a step 102 is performed of prompting a user to select a  
15 currently executing application program to be terminated. In the actual  
16 embodiment of the invention, the user is required to respond to this prompt; all  
17 other program activity ceases until the user responds.

18 Prompting the user is accomplished by way of a “System Out of Memory  
19 Dialog.” This dialog box is a special, “system modal” dialog box that essentially  
20 freezes the rest of the system. The user is informed that memory is critically low,  
21 and is forced to choose which applications should be shut down.

22 Step 104 is then performed of terminating the selected program, without  
23 allowing its further execution. All of the selected program’s threads are destroyed  
24 immediately, and all resources used by the program are closed or freed. This  
25 potentially causes a loss of data. However, the threshold that causes this is only

1       encountered when memory usage has reached a point where the system might  
2       become unstable without immediate remedial action. Thus, the drastic step of  
3       terminating a program is justified.

4       If the result of comparison 100 is negative, the operating system performs a  
5       second test 106, comparing current memory usage or availability against a usage  
6       or availability threshold that is referred to as an “intermediate” threshold. This  
7       threshold, referred to as a low memory threshold, is less critical than the critical  
8       memory threshold described above, and is reached while there is still enough  
9       memory available so that any particular application program can safely be allowed  
10      to shut itself down. If the result of test 106 is positive, a step 108 is performed of  
11      prompting the user to select a currently executing application program to be closed  
12      (using the System Out of Memory Dialog described above). Again, the user is  
13      required to respond to this prompt; all other program activity ceases until the user  
14      responds.

15      Step 110 is then performed of requesting the application to close itself. In  
16      the Windows® environment, this is done by sending a WM\_CLOSE message to  
17      the application’s message loop, giving the same effect as if the user had pressed  
18      the application’s *Close* button. The operating system gives the application a  
19      predefined time in which to close itself, such as eight seconds. If after that  
20      predefined time, memory is still low, the dialog box will reappear. Also, if at any  
21      time during those eight seconds, more memory is requested such that the amount  
22      of free memory goes below the critical memory threshold, the “System Out of  
23      Memory Dialog” discussed in conjunction with step 102 will immediately appear.  
24      Accordingly, applications should be configured to shut down without allocating  
25      very much additional memory.

If the result of comparison 106 is negative, the operating system performs a third test 112, comparing current memory usage or availability against a usage or availability threshold that is referred to as the "hibernation" critical threshold. This threshold is less critical than the other two thresholds, and thus occurs when there is still a relatively higher amount of available or free physical memory. If the result of test 112 is positive, a step 114 is performed of sending a "hibernate" message from the operating system to at least one of the application programs requesting it to minimize its current use of memory.

In the Windows® environment, only one of the concurrently executing programs is "active" at any particular time, even though multiple programs might be executing and partially visible. The active program is typically the one that has the focus, and is the one to which keyboard and mouse inputs are directed. In the Windows® CE environment, specifically, currently executing application programs are listed on the taskbar which is displayed to the user. The operating system maintains a "Z" order, indicating the relative positions of the applications' windows from front to back. This order also indicates the order in which the applications were last active: the rearmost application program is the least recently active. In the actual embodiment of the invention, the hibernate message of step 114 is sent to a particular application program that has been least recently active. Thus, a currently active program will not be affected, at least initially. In practice, such messages are sent to a plurality of programs, in order according to which programs were least recently active. The messages are sent with short delays between messages, until available memory is again above the hibernation threshold.

1           The particular mechanism for sending a message to an application program  
2 will vary depending on the particular operating system or computer architecture in  
3 use. In the Windows® environment, however, each application program has a  
4 main window that implements what is commonly referred to as a message loop.  
5 The operating system sends different types of messages to the application  
6 programs' message loops, such as messages that indicate keyboard or mouse  
7 activity. In the preferred embodiment of the invention, the message of step 114 is  
8 sent from the operating system through an application's message loop as a  
9 message having a pre-defined message ID.

10          An application program can respond to a hibernation message in any way it  
11 sees fit. Non-compliant application programs might simply ignore the messages.  
12 Ideally, however, an application program will take steps to minimize or reduce its  
13 current use of virtual and physical memory. Preferably, an application will first  
14 determine whether any portion of the application is visible to the user. If it is not  
15 visible, the application should store all data structures representing current context  
16 or state information, and release all other resources that can be reconstructed if and  
17 when needed in the future. Thus, the user will perceive no unusual activity, except  
18 perhaps a slight slowing of the system.

19          Specifically, in response to a hibernate message an application should  
20 perform the following steps:

21           (1)     Free any large pieces of memory (such as caches) that were  
22                previously allocated by the operating systems virtual memory  
23                allocation function (VirtualAlloc in Windows® CE).

- 1                   (2) Free as many GWES (graphics/window/event manager) objects as
- 2                   possible. Such objects include windows, bitmaps, and device
- 3                   contexts.
- 4                   (3) Save state information and data to persistent storage for later
- 5                   restoration, and then free as much of the application heap as
- 6                   possible.

7                   Although Fig. 3 is useful for understanding the invention conceptually, the  
8                   comparison steps are not necessarily performed in a sequence as suggested in the  
9                   flowchart. In operation, the hibernation memory threshold is the first line of  
10                  defense against low memory situations. Step 112, which checks available memory  
11                  against the hibernation threshold, is actually performed at periodic intervals such  
12                  as every five seconds. If the hibernation memory usage threshold has been  
13                  reached at any particular time, a hibernate message is sent to the least recently  
14                  active application that has not already received such a message. The operating  
15                  system maintains a flag indicating which applications have been sent such  
16                  messages. If the condition continues, hibernation messages will be sent to  
17                  additional applications.

18                  No explicit action is taken to “unhibernate” an application. Rather, the  
19                  application program reactivates itself as a result of activation by the user through  
20                  normal Windows® CE commands such as by clicking on the program’s icon on  
21                  the taskbar. The operating system notes when an application program is made  
22                  active and resets the flag to indicate that the program is no longer in a hibernation  
23                  state.

24                  In addition to “hibernating” applications during periods of excessive  
25                  memory usage, the operating system takes further memory conservation measures

when the hibernation threshold has been reached. For instance, the operating system shell will not allow browsing the file system using new browser windows (although existing browser windows are allowed to function normally). In addition, the operating system shell will not allow its internal icon display cache to be modified. Under normal memory conditions, the operating system shell will purge and rebuild its internal icon display cache periodically. In a hibernating state the shell needs to maintain consistency, and might not be able to rebuild its cache.

In spite of hibernation messages, the system may still run critically low on memory. This can happen, for example, when an active application continues to demand more memory while non-active applications are being hibernated. The two remaining memory thresholds (the low memory threshold and the critical memory threshold) are encountered when the system is very close to failing a memory allocation call. If this happens, the system displays the memory handler dialog box as described above. The user is informed that memory is critically low, and is forced to choose which application or applications should be shut down. After the dialog box is dismissed, each selected application is closed as described above. As a further memory conservation measure, the operating system shell will not allow a user to launch new applications from the operating system's user interface (such as from the "run" dialog box or by double clicking icons) when available memory is below the intermediate threshold.

The low memory threshold and critical memory threshold are checked (steps 100 and 106) when application programs attempt to allocate more virtual memory by invoking the *VirtualAlloc* interface of the Windows® CE operating

1 system. The operating system maintains the following precautions in helping to  
2 prevent low memory conditions from occurring:

3 • When an application allocates memory, the system “filters” this  
4 request by limiting the memory size. This prevents a single application  
5 from stealing all available memory with one large allocation. When the  
6 system enters a low memory situation, the maximum memory limit is  
7 further reduced.

8 • The system monitors available memory and handles two levels of  
9 memory conditions: the Low Memory threshold, and the Critical Memory  
10 threshold (in addition to the hibernation threshold discussed above).

11 The table below defines values for these thresholds and maximum memory  
12 allocations for the two levels of low memory conditions, as used in the present  
13 embodiment of the invention. The names are defined simply to aid in describing  
14 the low memory scenarios below. The actual values are examples which might  
15 vary in different systems and with different memory configurations.

16 <b>Item</b>	17 <b>Value</b>	18 <b>Description</b>
19 <b>HIBERNATION_THRESHOLD</b>	20      128K	21      Minimum available system 22      memory, below which defines a 23      “low memory” system. A 24      WM_HIBERNATE message is sent 25      when the system memory falls below this value.
26 <b>LOWMEM_THRESHOLD</b>	27      64K	28      Minimum available memory size 29      system must maintain when system 30      is in a “low” memory state.
31 <b>LOWMEM_MAXSIZE</b>	32      16K	33      Maximum size allowed to be 34      allocated (using VirtualAlloc) when 35      system is in a “low” memory state.
36 <b>CRITMEM_THRESHOLD</b>	37      16K	38      Minimum available memory size 39      system must maintain when system 40      is in a “critical” memory state.
41 <b>CRITMEM_MAXSIZE</b>	42      8K	43      Maximum size allowed to be 44      allocated (using VirtualAlloc) when

system is in a “critical” memory state.

Under a low memory condition, the system's low memory handler appropriately responds to basically four situations:

- An application calls *VirtualAlloc* requesting a memory size greater than *LOWMEM\_MAXSIZE*. Any *VirtualAlloc* requesting a memory size greater than *LOWMEM\_MAXSIZE* will be failed if it would cause the amount of free physical memory in the system to go below the *LOWMEM\_THRESHOLD*. No System Out of Memory Dialog will be displayed.
- An application calls *VirtualAlloc* requesting a memory size less than *LOWMEM\_MAXSIZE*. If a *VirtualAlloc* requesting a memory size less than *LOWMEM\_MAXSIZE* would cause the amount of free physical memory in the system to go below the *LOWMEM\_THRESHOLD*, then the System Out of Memory Dialog is displayed. The user is allowed to either select some applications that the system will try to close, or to get more memory by taking it from the file system. The system sends WM\_CLOSE messages to the applications selected by the user. If an application marked to be closed does not shutdown within 8 seconds, an *End Task/Wait* dialog is displayed, giving the user the choice of terminating the application or waiting some more time.
- An application calls *VirtualAlloc* requesting a memory size greater than *CRITMEM\_MAXSIZE*. Any *VirtualAlloc* requesting a memory size greater than *CRITMEM\_MAXSIZE* will be failed if it would cause the amount of free physical memory in the system to go below the *CRITMEM\_THRESHOLD*.

1      • An application calls *VirtualAlloc* requesting a memory size less than  
2      *CRITMEM\_MAXSIZE*. If a *VirtualAlloc* requesting a memory size less  
3      than *CRITMEM\_MAXSIZE* would cause the amount of free physical  
4      memory in the system to go below the *CRITMEM\_THRESHOLD*, then the  
5      System Out of Memory Dialog is displayed. The user is allowed to either  
6      select some applications that the system will try to close, or to get more  
7      memory by taking it from the file system. The system calls  
8      *TerminateProcess* to terminate any applications selected by the user.

9      Before a call to *VirtualAlloc* is allowed to fail, additional steps are taken  
10     within operating system 44 to minimize physical memory usage. For example,  
11     stack scavenging is employed periodically to recover unused stack space. Unused  
12     stack space is virtual memory that has been committed for use by a stack, but that  
13     is not currently being used. Unused stack space results when a stack grows and  
14     then shrinks. When entire unused pages are left beyond the current stack pointer,  
15     such pages can be discarded or reclaimed for other use. Page discarding is another  
16     tactic that is used by operating system 44 to minimize the use of memory before  
17     failing a call to *VirtualAlloc*. Using this tactic, read-only pages are discarded if  
18     they have not been used recently.

19     In operating system 44, steps of reclaiming unused stack memory and  
20     discarding read-only memory pages are also performed as a result of reaching  
21     defined memory thresholds. These thresholds are preferably set in relation to the  
22     three thresholds described above, so that the reclaiming and discarding steps occur  
23     at slightly less critical thresholds than the more active steps shown in Fig. 3. In  
24     other words, reclaiming and discarding steps occur before requesting an  
25

1 application to minimize its memory usage, before requesting an application  
2 program to close itself, and before terminating an application program.

3 Alternatively, an operating system in accordance with the invention might  
4 initiate reclaiming and discarding steps just prior to steps 100, 106, and 112 of Fig.  
5 3. If these steps were to free enough memory, further memory freeing steps could  
6 be skipped.

7 The invention provides an effective way of preventing memory usage from  
8 exceeding safe limits. The actions taken at less critical thresholds interfere as little  
9 as possible with a user's actual work, while increasing levels of memory usage  
10 result in necessarily more intrusive measures. However, these more intrusive  
11 measures are often avoided by the actions resulting from the less critical  
12 thresholds. Thus, a user is able to make very effective use of the limited physical  
13 memory available in popular, inexpensive computer devices.

14 In compliance with the statute, the invention has been described in  
15 language more or less specific as to structural and methodical features. It is to be  
16 understood, however, that the invention is not limited to the specific features  
17 described, since the means herein disclosed comprise preferred forms of putting  
18 the invention into effect. The invention is, therefore, claimed in any of its forms  
19 or modifications within the proper scope of the appended claims appropriately  
20 interpreted in accordance with the doctrine of equivalents.